

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant : Arieh Don et al. Art Unit : 2163
Serial No. : 10/749,692 Examiner : Marcin R. Filipczyk
Filed : December 29, 2003 Conf. No. : 2149
Title : DATA VERIFICATION FOLLOWING DATABASE WRITE

Mail Stop Appeal Brief - Patents

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

BRIEF ON APPEAL

(1) Real Party in Interest

The real party in interest is EMC Corporation, a corporation of Massachusetts having a place of business at 176 South Street, Hopkinton, Massachusetts 01748, as evidenced by assignments executed between December 16, 2003 and December 19, 2003, and recorded at the U.S. Patent Office on May 19, 2005 at Reel 016259, Frame 0488.

(2) Related Appeals and Interferences

There are no related appeals or interference.

(3) Status of Claims

Claims 1-17 are pending and on appeal. Of these, claims 1, 9, 14, 16, and 17 are independent.

(4) Status of Amendment

An amendment under Rule 41.33(a) is being filed with this paper to address the examiner's 101 rejections of claims 13-14, and objections to claims 9-13, and 17. The amendment has not been entered.

(5) Summary of Claimed Subject Matter

1. In a data storage system that stores both database records and ancillary data, a method	Ancillary data includes, for example, "data in temporary swap files, log data, and index
---	--

for processing a request from a host to write a database record to a target location on a logical device associated with a data-storage system in data communication with the host, the logical device having a first set of extents designated for storage of database records and a second set of extents designated for storage of ancillary data, the method comprising:	data.” (page 5, lines 20-21) The “target location” is where one would expect the data to be written. (step 58 in FIG. 7, and page 8, lines 5-11) In an implementation described in the application, the “logical device having a first set of extents designated for storage of database records and a second set of extents designated for storage of ancillary data.” is formed, for example, from combining physical devices, or portions thereof, on which data is stored.” (page 4, lines 21-23) The logical device includes, for example, a “user data section 22,” which is further divided into “data stores 32.” (FIG. 2, and page 5, line 10-14) These data stores 32 are referred to as “extents.” (page 5, lines 10-11)
maintaining, at the data storage system, information for distinguishing between the first set of extents and the second set of extents; based on the information, determining	In an implementation described in the application, the “information for distinguishing between the first set of extents and the second set of extents” include, for example, an “extent table 20.” (page 5, line 25 to page 6, line 20). The step of maintaining this information at the data storage system is disclosed on, for example, page 4, line 23.
that the target location is on an extent that is designated for storage of a database record, and	The step of “determining that the target location is on an extent that is designated for storage of a database record” is described in, for example, step 58 of FIG. 7, and on page 8, lines 5-11.
that the target location is not on an extent that is designated for storage of ancillary data	The step of “[determining] that the target location is not an extent that is designated for storage of ancillary data” is described in, for example, page 8, lines 5-11.
whereby the database record is not written to an extent that is designated for storage of ancillary data; and	In one implementation described in the application, if the extent is “designated for storage of ancillary data,” then the write is improper, and a “data storage system 14 returns an error message.” (page 8, lines 12-14, and steps 60 and 62 in FIG. 7)

writing the database record only to an extent that is designated for storage of a database record.	In one implementation described in the application, if the extent is "designated for storage of a database record," then the write is deemed to be proper, and the "data storage system 14 writes the data to the appropriate target location. (page 8, lines 14-19, and step 68 in FIG. 7)
9. A method of processing an Input/Output request to access a storage device having a plurality of extents defined thereon, each of the extents having a corresponding set of processing instructions associated therewith, the method comprising:	<p>The Input/Output request includes one or more processing instructions or steps, for accessing a storage device, e.g., logical device 18. (page 10, line 18, to page 11, line 2)</p> <p>In an implementation described in the application, the storage device includes, for example, a "user data section 22," which is further divided into "data stores 32." (FIG. 2, and page 5, line 10-14) These data stores 32 are referred to as "extents." (page 5, lines 10-11)</p>
receiving an Input/Output request having an associated target location on the storage device;	The step of "receiving an Input/Output request having an associated target location on the storage device," is described in, for example, page 5, lines 25-27.
identifying an extent set associated with the target location, the extent set having at least a first extent and a second extent, each of the first and second extents having its own associated processing instructions;	The step of "identifying an extent set associated with the target location, the extent set having at least a first extent and a second extent, each of the first and second extents having its own associated processing instructions," is described in, for example, page 6, lines 3-13. (See also page 6, line 25, to page 7, line 6)
determining that execution of the processing instructions for a first extent does not preclude execution of processing instructions for the second extent;	The step of "determining that execution of the processing instructions for a first extent does not preclude execution of processing instructions for the second extent," is described in, for example, page 7, lines 3-16.
executing the Input/Output request;	The step of "executing the Input/Output request," is described in, for example, FIG. 7, steps 52-72 , and page 7, line 25, to page 8, line

	27.
<p>executing the processing instructions associated with the first extent; and</p> <p>executing the processing instructions associated with the second extent.</p>	<p>The steps of "executing the processing instructions associated with the first extent," and "executing the processing instructions associated with the second extent," is described in, for example, page 6, lines 3-11.</p>
<p>14. A data-storage system comprising:</p> <p>a logical device having a plurality of extents defined thereon, each of the extents having a corresponding set of processing instructions associated therewith,</p>	<p>In an implementation described in the application, the "logical device having a plurality of extents defined thereon," is formed, for example, from combining physical devices, or portions thereof, on which data is stored." (page 4, lines 21-23) The logical device includes, for example, a "user data section 22," which is further divided into "data stores 32," (FIG. 2, and page 5, line 10-14) These data stores 32 are referred to as "extents." (page 5, lines 10-11).</p> <p>Each of the extents have "a corresponding set of processing instructions associated therewith," for example, each extent is associated with "a set of verification flags indicating what data verification tests, if any, are to be performed on any data written to the extent." (page 6, lines 6-9)</p>
<p>the logical device having</p> <p>a first extent having an associated first set of processing instructions, and</p> <p>a second extent having an associated second set of processing instructions;</p> <p>wherein the first set of processing instructions includes instructions that are different from instructions in the second set of processing instructions; and</p>	<p>A logical device having "a first extent having an associated first set of processing instructions," and "a second extent having an associated second set of processing instructions," in which, "the first set of processing instructions includes instructions that are different from instructions in the second set of processing instructions," is described in for example, page 6, lines 3-13.</p>
<p>information identifying each extent on the logical device and the processing instructions associated with that extent.</p>	<p>In an implementation described in the application, "information identifying each extent on the logical device and the processing</p>

	<p>instructions associated with that extent," includes, for example, an "extent-table 20," each extent-table entry 36 [including] locating information 38, ... and a set of verification flags 39 indicating what data verification tests, if any, are to be performed on any data written to [an] extent." (page 6, lines 3-11)</p>
<p>16. A computer-readable medium having encoded thereon software for causing a data storage system that stores both database records and data other than database records to process a request from a host to write a database record to a target location on a logical device associated with a data-storage system in data communication with the host, the logical device having a first set of extents designated for storage of database records and a second set of extents designated for storage of data other than database records, the software comprising instructions that, when executed, cause a computer to:</p>	<p>In an implementation described in the application, "data other than database records" includes, for example, "ancillary data." (page 5, paragraph 4) Ancillary data includes, for example, "data in temporary swap files, log data, and index data." (page 5, lines 20-21) The "target location" is where one would expect the data to be written. (step 58 in FIG. 7, and page 8, lines 5-11)</p> <p>In an implementation described in the application, the "logical device having a first set of extents designated for storage of database records and a second set of extents designated for storage of data other than database records," is formed, for example, from combining physical devices, or portions thereof, on which data is stored." (page 4, lines 21-23) The logical device includes, for example, a "user data section 22," which is further divided into "data stores 32." (FIG. 2, and page 5, line 10-14) These data stores 32 are referred to as "extents." (page 5, lines 10-11)</p> <p>In an implementation described in the application, "[the] computer-readable medium having encoded thereon software for causing a data storage system that stores both database records and data other than database records to process a request from a host to write a database record to a target location on a logical device associated with a data-storage system in data communication with the host," includes, for example, the "database application 12," which communicates with "data storage system 14." (page 4, lines 17-20)</p>

maintain, at the data storage system, information for distinguishing between the first set of extents and the second set of extents; and	In an implementation described in the application, the "information for distinguishing between the first set of extents and the second set of extents" include, for example, an "extent table 20." (page 5, paragraph 5 to page 6, paragraph 3)
based on the information, determine that the target location is on an extent that is designated for storage of a database record, and	The step of "[determining] that the target location is on an extent that is designated for storage of a database record," is described in, for example, step 58 of FIG. 7, and on page 8, lines 5-11.
that the target location is not on an extent that is designated for storage of ancillary data	The step of "[determining] that the target location is not an extent that is designated for storage of ancillary data" is described in, for example, page 8, lines 5-11.
whereby the database record is written only to an extent that is designated for storage of a database record, and whereby the database record is not written to an extent that is designated for storage of ancillary data.	In one implementation described in the application, if the extent is "designated for storage of a database record," then the write is deemed to be proper, and the "data storage system 14 writes the data to the appropriate target location. (page 8, lines 14-19, and step 68 in FIG. 7) In one implementation described in the application, if the extent is "designated for storage of ancillary data," then the write is improper, and a "data storage system 14 returns an error message." (page 8, lines 12-14, and steps 60 and 62 in FIG. 7)
17. A computer-readable medium having encoded thereon software for processing an Input/Output request to access a storage device having a plurality of extents defined thereon, each of the extents having a corresponding set of processing instructions associated therewith, the software including instructions that, when executed, cause a computer to:	The "computer-readable medium having encoded thereon software for processing an Input/Output request to access a storage device having a plurality of extents defined thereon," includes for example, the "database application 12," which communicates with "data storage system 14." (page 4, lines 17-20) The Input/Output request includes one or more processing instructions or steps, for accessing a

	<p>storage device, e.g., logical device 18. (page 10, line 18, to page 11, line 2)</p> <p>In an implementation described in the application, the storage device includes, for example, a "user data section 22," which is further divided into "data stores 32." (FIG. 2, and page 5, line 10-14) These data stores 32 are referred to as "extents." (page 5, lines 10-11)</p>
receive an Input/Output request having an associated target location on the storage device;	"[Receiving] an Input/Output request having an associated target location on the storage device," is described in, for example, page 5, lines 25-27.
identify an extent set associated with the target location, the extent set having at least a first extent and a second extent, each of the first and second extents having its own associated processing instructions;	"[Identifying] an extent set associated with the target location, the extent set having at least a first extent and a second extent, each of the first and second extents having its own associated processing instructions," is described in, for example, page 6, lines 3-13. (See also page 6, line 25, to page 7, line 6)
determine that execution of the processing instructions for a first extent does not preclude execution of processing instructions for the second extent;	"[Determining] that execution of the processing instructions for a first extent does not preclude execution of processing instructions for the second extent," is described in, for example, page 7, lines 3-16.
execute the Input/Output request;	"[Executing] the Input/Output request," is described in, for example, FIG. 7, steps 52-72 , and page 7, line 25, to page 8, line 27.
<p>execute the processing instructions associated with the first extent; and</p> <p>execute the processing instructions associated with the second extent.</p>	<p>"[Executing] the processing instructions associated with the first extent," and</p> <p>"[executing] the processing instructions associated with the second extent," is described in, for example, page 6, lines 3-11.</p>

CLAIM 2

Claim 2's additional limitation of "maintaining information for distinguishing between the first set of extents and the second set of extents comprises maintaining an extent table having extent table identifying properties, ... [the properties depending] on whether the extent is designated for storage of a database record and not designated for storage of ancillary data," is described in, for example, page 5, line 25 to page 6, line 20.

CLAIM 3

Claim 3's additional limitation of "selecting the properties to include information identifying a set of data verification steps, ... wherein the data verification steps depend on whether the extent is designated for storage of database records or for storage of ancillary data," is described in, for example, page 6, lines 3-13.

CLAIM 4

Claim 4's additional limitation of "identifying the logical device to be a logical device designated for storage of database records and not a logical device designated for storage of ancillary data" is described in, for example, page 5, lines 3-9, and page 6, lines 1-24.

CLAIM 5

Claim 5's additional limitation of "identifying a set of data verification steps to be carried out in connection with writing data to an extent, ... the set of data verification steps [depending] on whether the extent is designated for storage of database records or for storage of ancillary data," is described in, for example, page 6, lines 1-2, and lines 9-13.

CLAIM 6

Claim 6's additional limitation of "carrying out the data verification steps, ... the data verification steps [depending] on whether [an] extent is designated for storage of database records, or for storage of ancillary data," is described in, for example, page 6, lines 1-2, and lines 9-13 (See also, FIG. 7 steps 60-72, and associated text in page 8, lines 12-22.)

CLAIM 7

Claim 7's additional limitation of "determining that [a] target location is contained completely within an extent, wherein a target location that is not contained completely within an extent is indicative of a target location being on an extent that is designated for storage of ancillary data" is described in, for example, page 3, lines 7-11, and page 9, lines 1, to page 10, line 12.

CLAIM 8

Claim 8's additional limitation of "determining that the target location is contained completely within one or more extents, all of which share the same data verification steps," is described in, for example, page 8, lines 23-27, and page 9, lines 5-6, lines 9-10, and lines 16-17.

CLAIM 10

Claim 10's additional limitation that "receiving an Input/Output request comprises receiving a write request" is disclosed in, for example, page 5, lines 24-26, and page 7, lines 25-27.

CLAIM 11

Claim 11's additional limitation of "determining that none of the extents associated with the target location overlap with each other," is disclosed in, for example, page 8, lines 17-22.

CLAIM 12

Claim 12's additional limitation of "determining that the target location is contained completely within one or more extents" is disclosed in, for example, page 9, lines 9-10, and page 10, lines 1-2.

CLAIM 13

Claim 13's additional limitation of "determining that the first and second extents overlap, and that the processing instructions associated with the overlapping first and second extents are compatible," is described in, for example, page 7, lines 3-16.

CLAIM 15

Claim 15's additional limitation in which the "information identifying each extent comprises an extent table having an extent table entry corresponding to an extent on the logical device" is described in, for example, page 6, lines 3-4.

(6) Grounds of Rejection to be Reviewed on Appeal

- (i) Claims 1-17 are rejected under 35 U.S.C. 102(e) as being anticipated by *Furuike* (U.S. Patent No. 6,947,944).
- (ii) Claims 14 and 15 are rejected under 35 U.S.C. 101 because the claimed invention is directed to non-statutory subject matter.

(7) Argument

Law of Anticipation

"It is well settled that anticipation under 35 U.S.C. §102 requires the presence in a single reference of all of the elements of a claimed invention." *Ex parte Chopra*, 229 U.S.P.Q. 230, 231 (BPA&I 1985) and cases cited.

"Anticipation requires the presence in a single prior art disclosure of all elements of a claimed invention arranged as in the claim." *Connell v. Sears, Roebuck & Co.*, 220 U.S.P.Q. 193, 198 (Fed. Cir. 1983).

"This court has repeatedly stated that the defense of lack of novelty (i.e., 'anticipation') can only be established by a single prior art reference which discloses each and every element of the claimed invention." *Structural Rubber Prod. Co. v. Park Rubber Co.*, 223 U.S.P.Q. 1264, 1270 (Fed. Cir. 1984), citing five prior Federal Circuit decisions since 1983 including *Connell*.

In a later analogous case the Court of Appeals for the Federal Circuit again applied this rule in reversing a denial of a motion for judgment n.o.v. after a jury finding that claims were anticipated. *Jamesbury Corp. v. Litton Industrial Prod., Inc.*, 225 U.S.P.Q. 253 (Fed. Cir. 1985).

After quoting from *Connell*, "Anticipation requires the presence in a single prior art disclosure of all elements of a claimed invention arranged as in the claim," 225 U.S.P.Q. at 256, the court observed that the patentee accomplished a constant tight contact in a ball valve by a lip on the seal or ring which interferes with the placement of the ball. The lip protruded into the area where the ball will be placed and was thus deflected after the ball was assembled into the valve. Because of this constant pressure, the patented valve was described as providing a particularly good seal when regulating a low pressure stream. The court quoted with approval from a 1967 Court of Claims decision adopting the opinion of then Commissioner and later Judge Donald E. Lane:

[T]he term "engaging the ball" recited in claims 7 and 8 means that the lip contacts the ball with sufficient force to provide a fluid tight seal **** The Saunders flange or lip only sealingly engages the ball 1 on the upstream side when the fluid pressure forces the lip against the ball and never sealingly engages the ball on the downstream side because there is no fluid pressure there to force the lip against the ball. The Saunders sealing ring provides a compression type of seal which depends upon the ball pressing into the material of the ring. *** The seal of Saunders depends primarily on the contact between the ball and the body of the sealing ring, and the flange or lip sealingly contacts the ball on the upstream side when the fluid pressure increases. 225 U.S.P.Q. at 258.

Relying on *Jamesbury*, the ITC said, "Anticipation requires looking at a reference, and comparing the disclosure of the reference with the claims of the patent in suit. A claimed device is anticipated if a single prior art reference discloses all the elements of the claimed invention as arranged in the claim." *In re Certain Floppy Disk Drives and Components Thereof*, 227 U.S.P.Q. 982, 985 (U.S. ITC 1985).

Furuike's disclosure

At a broad level, *Furuike* addresses the problem of what to do when one wishes to inspect data in a compressed database record.

According to the prior art, one must first decompress the entire database record, even if the data to be inspected is only a small fraction of the entire database record. *Furuike* recognized that this adversely affected performance.

Furuike improved on the prior art by providing a way to leave a portion of a database record uncompressed so that one could more readily inspect the data in that portion, without the need to decompress the entire database record.

More specifically, *Furuike* describes a database apparatus for managing a database that includes records, each record having information in the form of "attributions," e.g., a name, telephone number, address, and position data (longitude, latitude, altitude) of a shop or company. When a search is carried out, a search key (e.g., telephone number of a shop) is received, and used to locate a record or an attribution corresponding to the record.¹

Because of limitations in storage capacity, the amount of data that can be stored in a database is limited. Accordingly, database records are generally stored in a storage medium only after having been compressed. When the records are read from the database, they are restored (decompressed) to their earlier condition. *Furuike* draws attention to a disadvantage of conventional database management systems in which a record is compressed as a single unit. In such systems, during a database search, the entire record needs to be decompressed before the database system can search that record or the record's attributions.²

To overcome the foregoing disadvantage, *Furuike* teaches a database management system that permits one to compress only part of a record, thus leaving a decompressed portion available for use in searching for a record's attributions. The disclosed system classifies data according to attributions and makes "attribution record groups" corresponding to each

¹ *Furuike*, col. 1, lines 25-35

² *Furuike*, col. 1, lines 37-46

attribution. Each attribution record group is then grouped together into a single unit, and the unit compressed to form a "database file."³ A "database definition file" includes definition data about the database file, i.e., "data indicating the types of attributions, or data regarding data compression of each data record having [an] individual attribution."⁴ A "data compression portion" performs the data compression process in accordance with the definition data.⁵

Furuike also describes a controller 7, that includes a database retrieving apparatus for retrieving a target record from a database file having multiple attribution record groups (A-E).⁶ *Furuike* describes the retrieval procedure as follows:

[A] data base retrieving apparatus retrieves a target record to be searched from a database file, which is made up of plural attribution record groups (A-E), each of which is compressed in a unit of each of the attribution record groups. Data decompressing means decompresses a particular attribution record group, which is to be searched, when a search request for searching the database file is received. Searching means searches for a target record containing a search key in the particular attribution record group. The data decompressing means further decompresses the other attribution record groups, which are different from the particular attribution record group, when the searching means finds the target record. As a result, the database managing apparatus can unnecessary steps for decompressing record data belonging to other attributions than the attribution to be searched for. Furthermore, it can retrieve the requested record in a short time.⁷

As described above, *Furuike* teaches a database management apparatus for storing and retrieving compressed records by dividing the records into attribution record groups (A-E) and compressing or decompressing only a relevant attribution record group.

Furuike further teaches that the search key typically corresponds to attribution record group A, which is frequently searched as a result.⁸ Accordingly, the controller 7 does not compress the data in attribution record group A, but compresses the data in attribution record groups B-E.

³ *Furuike*, col. 2, lines 10-23

⁴ *Furuike*, col. 3, lines 52-60

⁵ *Furuike*, col. 3, lines 56-60.

⁶ *Furuike*, col. 2, lines 24-27

⁷ *Furuike*, col. 2, lines 24-41

⁸ *Furuike*, col. 5, lines 9-20; see also Abstract, lines 9-15.

Although *Furuike* refers to classifying database records based on “attributions,” these attributions have nothing to do with choosing the particular extents in which one would write the data. The *Furuike* attributions have nothing to do with classifying data as either a database record, or as ancillary data. *Furuike*’s attributions simply identify which portions of a database record are to be compressed and which portions of the database record remain uncompressed. This permits selected data from a database record to be inspected without having to decompress the entire database record.

Applicant’s disclosure

Applicant recognized that there are two kinds of data: the important data, and the less-important ancillary data. Applicant further recognized that there was no point in subjecting this ancillary data to data verification. Thus, the technical problem Applicant faced was how to avoid data verification of ancillary data while also ensuring data verification of database records.

Applicant’s solution was a method for processing a request from a host to write a database record to a target location on a logical device associated with a data-storage system in data communication with the host. The method includes maintaining, at the data storage system, information identifying those extents of the logical device that are specifically designated only for storage of database records and, on the basis of the information, determining whether the target location is one on which a database record is permitted to be stored.

Applicant’s disclosure has nothing to do with compressing or decompressing data based on classification of data in database records into attributions. In fact, Applicant’s disclosure has nothing to do with data compression at all.

The Examiner’s rejections are considered in more detail below.

Claims 1-17 are rejected under 35 U.S.C. 102(e) as being anticipated by *Furuike*

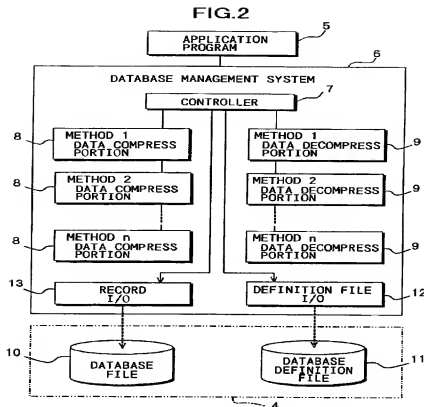
This section 102 rejection is improper for at least the following reasons:

1. Because *Furuike* fails to teach writing a database record to a target location on a “logical device having a first set of extents designated for storage of database records and a second set of extents designated for storage of ancillary data.”
2. Because *Furuike* fails to teach maintaining “information for distinguishing between the first set of extents and the second set of extents.”
3. Because *Furuike* fails to teach processing an Input/Output request having an associated target location on a storage device including identifying an extent set associated with the target location having “at least a first extent and a second extent, each of the first and second extents having its own associated processing instructions.”
4. Because *Furuike* fails to teach that maintaining information for distinguishing between first set of extents and second set of extents includes “maintaining an extent table having extent table entries identifying properties associated with the extent, wherein the properties depend at least in part on whether the extent is designated for storage of a database record and not designated for storage of ancillary data.”

Section 102 rejection of claims 1 and 16

***Furuike* fails to teach writing a database record to a target location on a “logical device having a first set of extents designated for storage of database records and a second set of extents designated for storage of ancillary data”**

In rejecting claim 1, the Examiner draws particular attention to FIG. 2, which is reproduced below.



The Examiner appears to regard claim 1's "logical device" as corresponding to the memory **4** in *Furuie*'s FIG. 2. The "first set of extents" in claim 1 would then correspond to the extents on which the database file **10** has been written. The "second set of extents" would then be the particular extents on which the database definition file **11** has been written. With this being the case, the Examiner appears to regard claim 1's "database record" as corresponding to *Furuie*'s data base file **10**, and claim 1's "ancillary data" as being the "database definition file **11**" in FIG. 2.

These correspondences are summarized below:

CLAIM 1 LIMITATION	ALLEGED DISCLOSURE IN <i>FURUIKE</i>
"logical device"	Memory 4 in FIG. 2
"database record"	Database file 10 in FIG. 2
"ancillary data"	Database Definition File 11 in FIG. 2

"first set of extents"	The particular extents on which database file 10 has been written.
"second set of extents"	The particular extents on which database definition file 11 has been written.

Having made the foregoing correspondences, the Examiner assumes that *Furuike's* controller **7** somehow maintains information that identifies which extents in memory **4** are reserved for a database file **10** and which extents are reserved for a database definition file **11**.

Applicant agrees that *Furuike's* controller **7** knows enough to store certain kinds of data in a database file **10** and to store certain other kinds of data in a database definition file **11**. Applicant also agrees that these two files are necessarily stored on different extents, since otherwise one would overwrite the other.

However, there is nothing in *Furuike* to suggest that the extents into which the controller **7** writes the database file **10** are specifically designated for database files, or that the extents on which the controller **7** writes the database definition file **11** are somehow specifically designated for database definition files.

It is true that once a database file **10** is written into a set of extents, that set of extents is currently in use by the database file **10**. However, this does not mean those extents were ever designated for database files. It is more likely that those extents were used because they happened to have been free at the time the controller **7** was looking for a place to write the database file **11**. Nothing would stop the controller **7** from later re-using those same extents to write a database definition file **11**.

In contrast, in Applicant's claim **1**, certain extents are designated for storage of ancillary data whereas other extents are designated for storage of database records. This is an important distinction.

Specifically, by designating extents solely for storage of ancillary data, Applicant's system recognizes that any data stored on such extents must be ancillary data. As a result, when

it is time to carry out data verification of stored data, Applicant's data verification software can bypass those extents, thus avoiding the waste of time associated with performing data verification on ancillary data.

In *Furuike's* system, all extents within the memory 4 are treated the same way. There is no indication that some extents are designated for ancillary data and some are designated for database records. As a result, there is no way for *Furuike* to avoid executing data verification on ancillary data.

Claim 16 recites a computer-readable medium having encoded thereon software for causing a data storage system to process a request from a host to write a database record to a target location on a "logical device having a first set of extents designated for storage of database records and a second set of extents designated for storage of ancillary data." For at least similar reason given above for claim 1, claim 16 is patentable over *Furuike*.

***Furuike* fails to teach maintaining "information for distinguishing between the first set of extents and the second set of extents."**

As best understood, the Examiner regards the controller 7 as maintaining information for distinguishing between a first set of extents, which is designated for storage of database records, and a second set of extents, which is designated for storage of ancillary data.

The controller 7 of *Furuike* performs a variety of tasks related to managing data compression, decompression and searching.⁹ For example, *Furuike* states:

A data controller (data compressing means, data decompressing means, searching means) 7 has a function for operating the input/output device 2, and controls a data compression portion (data compressing means) 8 (1, 2, . . . , n) for performing data compression processes (compression methods) and a data decompression portion (data decompressing means) 9 (1, 2, . . . , n) for performing data decompression processes (decompression methods).¹⁰

The foregoing recitation of controller functions has to do with controlling data compression. Nothing in the foregoing recitation of what the controller 7 does suggests that it somehow maintains information about where particular types of data, whether compressed or

⁹ *Furuike*, col. 3, lines 41-47

¹⁰ *Furuike*, col. 3, lines 41-48

otherwise, are permitted to be written. Specifically, nothing in *Furuike*'s list of controller functions suggests that the controller 7 maintains information about the extents on which ancillary data may be written and the extents on which database records may be written.

In an example implementation described in the Applicant's disclosure, the "information for distinguishing between the first set of extents and the second set of extents" refers to the extent-table 20. In this regard, the application states:

The extent-table 20 contains an extent-table entry 36 corresponding to each data store 32 on its associated logical device 18. In one implementation of a logical device 18, there can be as many as 512 data stores 32. Hence, the extent-table 20 for a particular logical device 18 can have as many as 512 extent-table entries 36. Each extent-table entry 36 includes locating information 38 indicating the start address and size of an extent, and a set of verification flags 39 indicating what data verification tests, if any, are to be performed on any data written to that extent.¹¹

Nothing in *Furuike* suggests that the controller 7 maintains any information like the extent-table 20 "for distinguishing between the first set of extents and the second set of extents" as recited in Applicant's claim 1.

For at least the reasons above, controller 7 cannot be said to maintain "information for distinguishing between the first set of extents and the second set of extents."

Claim 16 recites, in part, a computer-readable medium having encoded thereon software, the software comprising instructions that, when executed, cause a computer to "maintain ... information for distinguishing between the first set of extents and the second set of extents." For at least a similar reason given for claim 1 above, claim 16 is patentable over *Furuike*.

Section 102 rejection of claims 9, 14, and 17

***Furuike* fails to teach processing an Input/Output request having an associated target location on a storage device including identifying an extent set associated with the target location having "at least a first extent and a second extent, each of the first and second extents having its own associated processing instructions."**

As best understood, the Examiner appears to regard claim 9's "storage device" as corresponding to the memory 4 in *Furuike*'s FIG.2. Accordingly, the "first extent" in claim 9

¹¹ *Specification*, page 6, lines 3-13

would then correspond to an extent on which the database file **10** has been written. The “second extent” would then be a particular extent on which the database definition file **11** has been written. *Furuike*’s database file **10** and database definition file **11** store data records and “definition data regarding the database file **10**,” respectively¹².

With this being the case, Applicant concedes that each of the data base file **10** and the database definition file **11** should then have its own associated processing instructions to provide guidance on how each of these items is to be used. However, these instructions are associated with *files*, not extents.

Nothing in *Furuike* suggests that the *extents* on which the database file **10** and the database definition file **11** are written would have their own associated processing instructions.

In contrast, in Applicant’s claim 9, each of first and second extents have their own associated processing instructions. These instructions are associated with the extents themselves. They apply to whatever data may be stored in those extents. This is an important distinction. Specifically, by associating, for example, data verification steps with a first extent, and bypassing the data verification steps for a second extent, Applicant’s system provides a way for data to be treated differently depending on where it is saved.

In *Furuike*, the processing instructions are associated with the type of data, e.g., whether the data is a database file **10** or a database definition file **11**, and not with the extents. As such, there is no indication anywhere in *Furuike* that some extents have different associated processing instructions from other extents.

Claim 14 recites a data storage system including a logical device having “a first extent having an associated first set of processing instructions, and a second extent having an associated second set of processing instructions.” For at least a similar reason given for claim 9, claim 14 is patentable over *Furuike*.

Claim 17 recites a computer-readable medium having encoded thereon software including instructions that, when executed, cause a computer to, among other steps, receive an


¹² *Furuike*, col. 3, lines 49-54

Input/Output request having an associated target location on a storage device, and identify an extent set associated with the target location, the extent set having at least "a first extent and a second extent, each of the first and second extents having its own associated processing instructions." For at least a similar reason given for claim 9, claim 17 is also patentable over *Furuike*.

Section 102 rejection of claim 2

Furuike fails to teach that maintaining information for distinguishing between first set of extents and second set of extents includes "maintaining an extent table having extent table entries identifying properties associated with the extent, wherein the properties depend at least in part on whether the extent is designated for storage of a database record and not designated for storage of ancillary data."

The Examiner appears to regard the database definition file 11 to be claim 2's "extent table having extent table entries identifying properties associated with the extent." The role of this database definition file 11 is described in steps A1-5 of *Furuike*'s FIG. 3, reproduced below



According to *Furuike*, the database definition file **11** “includes several definition data regarding the database file **10**, that is, data indicating types of attributions, or data regarding data compression of each data record having individual attribution.”¹³ The system of *Furuike* uses the database definition file **11** (Steps A1-A3, of FIG. 3) to classify data according to attributions, and to make “attribution record groups” corresponding to each attribution (Steps A4-5). Each attribution record group is then grouped together into a single unit. The unit is then compressed to form a “database file.” (Steps A6-1 to A6-n, and A7-8) A “data compression portion” performs the data compression process in accordance with the definition data in the database definition file **11**. With this being the case, *Furuike*’s “attributions” simply identify which portions of a database record are to be compressed and which portions of the database record remain uncompressed.

This has nothing to do with claim 2’s limitation of

“maintaining an extent table having extent table entries identifying properties associated with the extent, wherein the properties depend at least in part on whether the extent is designated for storage of a database record and not designated for storage of ancillary data.”

In an example implementation described in the Applicant’s disclosure, the extent-table **20** is described as having extent-table entries **36** corresponding to each data store **32** on the logical device **18**. There can be as many as 512 data stores **32**. As such, the extent-table **20** for the logical device **18** can have as many as 512 extent-table entries **36**. Each extent-table entry **36** includes locating information **38** that indicates a start address and size of an extent, and a set of verification flags **39** indicating what data verification tests, if any, are to be performed on any data written to that extent.

Although *Furuike* refers to classifying database records based on “attributions,” these attributions have nothing to do with choosing the particular extents in which one would write the data. As such, the attributions in *Furuike* have nothing to do with classifying data as either a database record, or as ancillary data. Accordingly, nothing in *Furuike* suggests maintaining anything like an extent-table **20** having

¹³ *Furuike*, col. 3, lines 50-54

“extent table entries identifying properties associated with the extent, wherein the properties depend at least in part on whether the extent is designated for storage of a database record and not designated for storage of ancillary data”

as recited in Applicant's claim 2.

Claims 14 and 15 are rejected under 35 U.S.C. 101 because the claimed invention is directed to non-statutory subject matter.

Claim 14 has been amended. As amended, claim 14 recites a data-storage system comprising “a logical device having a plurality of extents defined thereon ... the logical device having ... a first extent ...and a second extent; and a computer readable medium having information identifying each extent on the logical device and the processing instructions associated with that extent.” Claim 14 clearly recites statutory subject matter. The MPEP states:

“[F]unctional descriptive material” consists of data structures and computer programs which impart functionality when employed as a computer component.

...

When functional descriptive material is recorded on some computer-readable medium, it becomes structurally and functionally interrelated to the medium and will be statutory in most cases since use of technology permits the function of the descriptive material to be realized. Compare *In re Lowry*, 32 F.3d 1579, 1583-84, 32 USPQ2d 1031, 1035 (Fed. Cir. 1994)(discussing patentable weight of data structure limitations in the context of a statutory claim to a data structure stored on a computer readable medium that increases computer efficiency) and *In re Warmerdam*, 33 F.3d 1354, 1360-61, 31 USPQ2d 1754, 1759 (claim to computer having a specific data structure stored in memory held statutory product-by-process claim) with *Warmerdam*, 33 F.3d at 1361, 31 USPQ2d at 1760 (claim to a data structure *per se* held nonstatutory).¹⁴

In claim 14, a data storage system includes a “computer-readable medium” having “information identifying each extent on the logical device and the processing instructions associated with that extent.” A computer-readable medium having such information is statutory under *Lowry* because that information's functionality, which is “identifying each extent on the logical device and the processing instructions associated with that extent” is realized.

For at least the reason above, Applicants submit that the subject matter of claims 14 and 15 is statutory and respectfully request that the § 101 rejection be withdrawn.

¹⁴ MPEP § 2106.01.

SUMMARY

This appeal brief follows the Notice of Appeal filed on June 5, 2008.

The required brief fee of \$540 and extension of time fee of \$490 is being paid concurrently on the Electronic Filing System (EFS) by way of Deposit Account authorization, referencing Attorney Docket No. 07072-0159001. No other fees are believed to be due in connection with the filing of this appeal brief. However, to the extent fees are due, or if a refund is forthcoming, please adjust our deposit account.

Respectfully submitted,

Date: October 6, 2008



Faustino A. Lichaúco
Reg. No. 41,942

Fish & Richardson P.C.
225 Franklin Street
Boston, MA 02110
Telephone: (617) 542-5070
Facsimile: (617) 542-8906

APPENDIX OF CLAIMS (including amendments proposed under Rule 41.33(a))

1. In a data storage system that stores both database records and ancillary data, a method for processing a request from a host to write a database record to a target location on a logical device associated with a data-storage system in data communication with the host, the logical device having a first set of extents designated for storage of database records and a second set of extents designated for storage of ancillary data, the method comprising:

maintaining, at the data storage system, information for distinguishing between the first set of extents and the second set of extents; based on the information, determining

that the target location is on an extent that is designated for storage of a

database record, and

that the target location is not on an extent that is designated for storage of ancillary data

whereby the database record is not written to an extent that is designated for storage of

ancillary data; and

writing the database record only to an extent that is designated for storage of a database

record.
2. The method of claim 1, wherein maintaining information for distinguishing between the first set of extents and the second set of extents comprises maintaining an extent table having extent table entries identifying properties associated with the extent,

wherein the properties depend at least in part on whether the extent is designated for storage of a database record and not designated for storage of ancillary data.

3. The method of claim 2, further comprising selecting the properties to include information identifying a set of data verification steps to be carried out when data is written into the extent,

wherein the data verification steps depend on whether the extent is designated for storage of database records or for storage of ancillary data.

4. The method of claim 1, further comprising identifying the logical device to be a logical device designated for storage of database records and not a logical device designated for storage of ancillary data.

5. The method of claim 1, further comprising identifying a set of data verification steps to be carried out in connection with writing data to an extent,

wherein the set of data verification steps depends on whether the extent is designated for storage of database records or for storage of ancillary data.

6. The method of claim 5, further comprising carrying out the data verification steps,

wherein the data verification steps depend on whether the extent is designated for storage of database records or for storage of ancillary data.

7. The method of claim 1, wherein determining whether the target location is on an extent that is designated for storage of a database record and that the target location is not on an

extent that is designated for storage of ancillary data comprises determining that the target location is contained completely within an extent, wherein a target location that is not contained completely within an extent is indicative of a target location being on an extent that is designated for storage of ancillary data.

8. The method of claim 3, wherein determining that the target location is on an extent that is designated for storage of a database record, and that the target location is not on an extent that is designated for storage of ancillary data comprises determining that the target location is contained completely within one or more extents, all of which share the same data verification steps.

9. A method of processing an Input/Output request to access a storage device having a plurality of extents defined thereon, each of the extents having a corresponding set of processing instructions associated therewith, the method comprising:

receiving an Input/Output request having an associated target location on the storage device;

identifying an extent set associated with the target location, the extent set having at least a first extent and a second extent, each of the first and second extents having its own associated processing instructions;

determining that execution of the processing instructions for a first extent does not preclude execution of processing instructions for the second extent;

executing the Input/Output request;

executing the processing instructions associated with the first extent; and

executing the processing instructions associated with the second extent.

10. The method of claim 9, wherein receiving an Input/Output request comprises receiving a write request.
11. The method of claim 10, further comprising selecting the processing instructions to include instructions for verifying that the writing of the data to the target location was carried out successfully.
12. The method of claim 9, wherein determining that execution of the processing instructions for a first extent does not preclude execution of processing instructions for the second extent comprises determining that none of the extents associated with the target location overlap with each other.
13. The method of claim 9, wherein determining that execution of the processing instructions for a first extent does not preclude execution of processing instructions for the second extent comprises determining that the first and second extents overlap, and that the processing instructions associated with the overlapping first and second extents are compatible.
14. A data-storage system comprising:

a logical device having a plurality of extents defined thereon, each of the extents having

a corresponding set of processing instructions associated therewith,

the logical device having

a first extent having an associated first set of processing instructions, and

a second extent having an associated second set of processing instructions;

wherein the first set of processing instructions includes instructions that are
different from instructions in the second set of processing instructions;
and

a computer readable medium having information identifying each extent on the
logical device and the processing instructions associated with that extent.

15. The system of claim 14, wherein the information identifying each extent comprises an extent table having an extent table entry corresponding to an extent on the logical device.
16. A computer-readable medium having encoded thereon software for causing a data storage system that stores both database records and data other than database records to process a request from a host to write a database record to a target location on a logical device associated with a data-storage system in data communication with the host, the logical device having a first set of extents designated for storage of database records and a second set of extents designated for storage of data other than database records, the software comprising instructions that, when executed, cause a computer to:

maintain, at the data storage system, information for distinguishing between the first set of extents and the second set of extents; and

based on the information, determine

that the target location is on an extent that is designated for storage of a
database record, and

that the target location is not on an extent that is designated for storage of
ancillary data

whereby the database record is written only to an extent that is designated for storage of
a database record, and

whereby the database record is not written to an extent that is designated for storage of
ancillary data.

17. A computer-readable medium having encoded thereon software for processing an
Input/Output request to access a storage device having a plurality of extents defined
thereon, each of the extents having a corresponding set of processing instructions
associated therewith, the software including instructions that, when executed, cause a
computer to:

receive an Input/Output request having an associated target location on the storage
device;

identify an extent set associated with the target location, the extent set having at least a
first extent and a second extent, each of the first and second extents having its own
associated processing instructions;

determine that execution of the processing instructions for a first extent does not
preclude execution of processing instructions for the second extent;

execute the Input/Output request;

execute the processing instructions associated with the first extent; and

execute the processing instructions associated with the second extent.

Applicant : Arie Don et al.
Serial No. : 10/749,692
Filed : December 29, 2003
Page : 32 of 33

Attorney's Docket No.: 07072-0159001 / EMC 03-132

EVIDENCE APPENDIX

None

Applicant : Arie Don et al.
Serial No. : 10/749,692
Filed : December 29, 2003
Page : 33 of 33

Attorney's Docket No.: 07072-0159001 / EMC 03-132

RELATED PROCEEDINGS APPENDIX

None